Compositionality and other issues in the philosophy of mind and language

Jönsson, Martin; Brinck, Ingar

Compositionality and Other Issues
in the Philosophy of Mind and Language
An interview with Jerry Fodor
by
MARTIN L. JÖNSSON and INGAR BRINCK
University of Lund

JERRY FODOR has been working on problems about primarily the representational properties of mind and psychosemantics for more than four decades. One of his central concerns is the compositionality of concepts and meaning, another the conceptual nature of the human mind. Martin Jönsson and Ingar Brinck arranged to meet Fodor in Lund in March 2005 to discuss Fodor's views on these matters. The major parts of the conversation are printed below, starting with the nature of compositionality, continuing with the relation between concepts and categorisation, then turning to the conceptual capacities of animals. Finally Fodor expresses his thoughts about conceptual analysis and its role for philosophy today, *viz.* to construct and clarify theoretical concepts in close collaboration with the empirical sciences.

1. Compositionality

MJ: When did you start thinking about compositionality and what got you started?

JF: Back ten, fifteen years ago I got involved in discussions with people at Stanford who for one reason or another wanted to run a story about semantics in which the units of interpretation are something like sentences without internal structure. The standard argument against that has always been productivity – you have got infinitely many sentences that cannot all be primitive. But these guys didn't like that argument.

I suggested that you can get the same thing with a much more conservative set of premises by arguing from the *systematicity* of language to its compositionality. You get symmetries like aRb and bRa. The natural way to explain that kind of symmetry is to say that suppose that the paired structures are constructed out of the same lexical constituents. This is a

standard argument now and that sort of led into thinking about compositionality.

At the same time there was a lot of discussion in the literature about stereotype theories of concepts. I don't know who first observed that it is very hard to see how you can get stereotypes to compose. It certainly wasn't me. Then it occurred to me and to Ernie Lepore, with whom I have occasionally worked on this, that you could actually generalize this argument. That is, it turns out that if you take any notion of concept individuation or concept possession that involves epistemic constraints, it will violate compositionality. The epistemic constraints for stereotypes, for instance, involve the notion of subjective probability. The reason they violate compositionality is essentially that the epistemic constraints don't themselves compose.

So it was the confluence of those two ideas. What Ernie and I have come to believe over the years is that there probably is not any way of meeting the compositionality constraint for a theory of language or a theory of thought, unless you assume that the semantics is referential. Reference looks like the only thing that does compose. If this line of thought is correct, then it turns out to be a very powerful constraint

MJ: And, in your opinion, one that *only* a referential theory of meaning will meet?

JF: There may be some way of having a semantic theory which has both reference and sense that is compositional. But I have no idea of how to do it, and I have seen no proposals. It looks like a general form of argument against almost every proposal for a theory of meaning that people have been running either in philosophy or in psychology for the last fifty years. I am quite pleased with that actually, being rather anarchistic by temperament.

IB: Did you ever discuss the basic definition of compositionality, or whether there can be different kinds of compositionality?

JF: That is a fair question. I don't know of anybody who has a useful working definition of compositionality. The intuition is clear enough. You project the semantic and also the syntactic properties of complex expressions, or thoughts or whatever, from some inventory of primitives together with their syntax. As far as I understand it, that is the way the serious grammar has been working recently, and I think that more and more syntactic structures have projections from lexical structures. So that is fine. But how exactly one says what it is for a system to be composi-

tional, I actually don't know. It is very easy to trivialize the condition because you can say that it has to be a function. But functions are cheap.

MJ: Let us make a distinction between a weak and a strong interpretation of compositionality. According to the weak interpretation there is a function from the meaning of the parts of a complex expression and the syntactical rule by which they have been combined to the meaning of that complex expression. According to the strong interpretation the meaning of a complex expression is identical to the meaning of the parts of that expression structured in a way signified by the syntactic composition of the parts. First of all, do you think there are any good arguments in favour of the stronger kind of relation?

JF: Well, there is a supervenience relation anyway. Whether there is an identity relation I don't know. And I don't know how you should think about the supervenience. You might think about it in terms of a logical construction or something of that kind.

Another way to think about research in this area is just to give up on saying what compositionality is in general, and describe how the mapping works in specific languages or in thought. That is obviously a stronger claim than just that there is a mapping. That is what happens in syntax: There is no general very useful characterization because the abstract characterization is pretty weak. What you actually do is make substantive claims about what kinds of recursion languages actually exhibit. Maybe that is the right thing to do in semantics too.

There is a general tendency, especially among philosophers, that if you want to work on cows as it might be, the way you do it is that you get a general definition of 'cow' and then look at some cows and see how they fit. This seems to me not to be the way to do work on compositionality. The process doesn't involve giving some general definition of compositionality and then try to fit it to natural language. The natural research project would be to go in the other direction, given that the intuition that some sort of compositional process is going on seems to be so strong.

MJ: Does this mean that you don't think that there are any *general* arguments in favor of the strong kind of compositionality?

JF: The major argument for the strong kind, I would have thought, was the lack of alternatives, and we do know that for some kinds of cases you get perfectly clear relations between the semantics of complex expressions and the semantics of their constituents. Brown cows are brown and brown cows are cows and that is no accident.

MJ: But given that these kinds of cases are supposed to ground an argument in favor of strong compositionality, cannot cases like 'big ant' (where 'big' is not entailed) ground an argument against strong compositionality?

JF: I don't think so. The problem with big ants, and with rising temperatures, is not that we don't know how to do them, but that there are too many ways of doing them. And we don't know which one to choose. But, there is no reason to doubt that you can get compositional analyses, it is that they are several in competition. So, I don't know of any cases where you got something that clearly doesn't give you compositional treatment, but doesn't feel like an idiom.

MJ: But if the meaning of 'big' is part of the meaning of 'big ant' why is not there an entailment relation from 'big ant' to 'big'?

JF: Well, there is, but it is not like conjunction, it is not intersection. 'Big ant' entails 'big for an ant'. That is the fact that has to be captured. And that is a perfectly kosher entailment, there is nothing wrong with that, it is just that you cannot do it with set intersection.

MJ: You have pursued other lines of argument in favour of strong compositionality. For instance in 'Why compositionality won't go away' you argue that given that natural languages exhibit a weak kind of compositionality as well as reverse compositionality[1], assuming the strong kind of compositionality explains why natural languages exhibit the other two kinds of compositionality. And this would be a reason to assume the stronger kind of compositionality – Is this something you still find compelling?

JF: Yeah. I think that the reverse compositionality stuff is important, because it blocks what strikes me as an intrinsically unattractive account. This issue usually comes up when talking about possession conditions (possession of concepts). The way it is supposed to work is that there are possession conditions of primitive concepts that they don't transmit to complexes. That will allow you to live with the fact that possession conditions don't compose, that is really why people want to run that line. But then you have trouble with reverse compositionality, with explaining why it is that somebody that has the complex concept also has the constituents. So for polemical purposes it is important that it goes in both directions.

---

[1] '…Each constituent expression contributes the whole of it is meaning to its complex hosts' (Fodor and Lepore 2002, p. 59)

For purposes of actually doing the research what is important is just that the intuitions be preserved.

MJ: I agree. Actually, for polemical purposes I think that the strong kind of compositionality let's you extend the compositionality critique to more theories than those with epistemic clauses. For instance, it seems to me that classical possible world semantics (on an extensional reading), is incompatible with strong compositionality. On this approach, traditionally the meaning of sentences are understood in terms of functions from possible worlds to truth values, while singular terms are functions from possible worlds to individuals, and predicates are functions from possible worlds to extensions. But such a distribution is clearly in violation of strong compositionality, although it is clearly compatible with the weak kind of compositionality.

JF: It depends on how you do it. It could be. If you take the basic semantic relation to be between a sentence, roughly speaking, and its extension across worlds, then you got a problem with this. But it is not clear to me that you cannot adapt a possible world semantics to a theory of internal structure for sentences.

MJ: In a non-ad hoc way? So that the adaptation is motivated independently?

JF: Yes. But possible worlds semantics for all sorts of reasons seems to me a sort of dubious bet about how the semantics ought to go. There are obvious problems about intentionality that everybody is aware of. It is not clear whether you can get that kind of thing to work as a research project. You don't want all the numerical truths to mean the same thing, but they look like having the same extension. Necessary co-extension is very hard to deal with on the possible world kind of model.

MJ: Since you have employed compositionality in arguments against a range of theories, and you seem to be leaning towards a strong interpretation of compositionality, the question arises: To what extent do you think your arguments against these theories depend on that compositionality is interpreted in a strong way?

JF: I don't think that it matters. The geography in this area is something like this: All the standard theories of concepts have taken it for granted that the way to understand concepts is via their possession conditions. So if you tell me what it is to have a concept, then I'll tell you what the concept is. And it turns out what it is to have the concept is invariably understood in epistemic terms. It is being able to sort things, being able to draw inferences, stuff like that – the standard ways to unpack use theories of expressions.

You can give a general argument that these theories are not going to compose on any construal of compositionality, supervenience or weaker. You just cannot put them together because the epistemic properties themselves are not compositional. So I think that that kind of argument doesn't turn on how robust your notion of compositionality is.

MJ: It is ironic that someone like you who has been so successful in using compositionality to refute theories of meaning and of concepts is inclined to advance a referentially based semantics, since Frege, who put that kind of argument on the modern agenda[2] an be understood as using compositionality as an argument against referential theories of meaning. So the question then arises: How do you defend your own position against Frege's arguments?

JF: I don't think that the Frege arguments are any good. That is a fair question though. There are really two worries for a referential semanticist, there is that and there is the problem of how you get publicity, in the sense of how more than one person can have a given concept or how more than one time-slice of a person can have a given concept. If the Frege arguments are any good, then what they show is that you have to have at least a two factor semantics (you have to have senses and references), and then there is an open question what the relation is between the factors and so on. But, I think that the arguments have been overplayed. There are straightforward answers to the standard arguments in terms of what you do about the concepts WATER and H_2O: They have different possession conditions, you cannot have the concept H_2O unless you have the concept HYDROGEN but you can have the concept WATER without having the concept HYDROGEN! So the concepts have different structures. So that suffices to pin the failure of substitution on.

What you get down to is: Suppose you have pure primitive expressions, which are not distinguishable by their history or something pragmatic, and they are Millian[3] names in the sense that they are not even psychologically condensed descriptions – Then does substitution fail? The answer is: It is very hard to tell.

_____

[2] In his application of the compositionality principle in terms of the near equivalent substitutivity principle (Frege 1892). See Hodges (2001) for equivalence results.

[3] I.e. that they don't have any meaning over and above their referent.

The more you nibble away at this stuff until you work down to the really pure cases, the weaker the intuitions get. And in fact, my impression is that they eventually come into conflict with intuitions about translation. Because, we have a language where a guy has two names, Jim and Sam. You have another language in which he has two names, *a* and *b*, and they all function like Millian names. Then the intuition is, I think, that you can translate anyway you like. There is nothing at issue in whether it is Jim that is translated with *b* or Sam that is. So I think that when you really get down to the cases, when push comes to shove and your back's to the wall, what you find is that the intuitions are pretty weak, and they seem to run against intuitions about translation.

So it may be that when you, as it were, polished up the argument to get rid of the cases that are really not showing anything, the intuitions that Frege is resting on are weak enough so that one could shrug one's shoulders at them. I mean you are after all going to have phonological differences between names. And the Mates kind of cases[4], suggest that even that can be recruited to block substitution if you have sufficiently contrived contexts. So it is really hard to make clear that one have a fund of impeccable cases of non-substitution. Until someone shows me a knock-down case, I'm inclined to think that maybe Frege overplayed his hand. You cannot even raise the Frege problems until you get down to purely primitive expressions. And that takes away a lot of, as it were, sexier examples, the ones that you impress undergraduates with – the Morning star and the Evening star for example.

MJ: Since compositionality has been such a big part of your philosophical project, your 2001 article 'Language, Thought and Compositionality', in which you claim that language probably is not compositional (only thought is) came as a bit of a surprise. The plot thickens a year later, when you, in the geographical introduction to 'The Compositionality Papers' (where a lot of your articles on compositionality are collected), claim that language (as well as thought) really is compositional. So, which is it?

JF: Well, there is an exoteric and an esoteric doctrine. The context of the compositionality arguments for years and years was: Look, something has to be productive and systematic, that is clear, and whatever is systematic and productive has to be compositional. And the natural assumption was that language was compositional – the easy examples are from lan-

guage. If you were having arguments with behaviorists or association-ists or something it didn't matter whether you took it to be language or thought. And language is a bit more perspicuous, the data is relatively public.

But if one turns one's attention to the question 'Is this a theory of language or is it a theory of thought?', then you have to make those kinds of distinctions. My guess is that the right treatment would say that natural languages like English, or whatever, are not compositional. In fact my guess is that they don't have a semantics at all. They have a syntax, basically.

What has a semantics, on this view, is thought, and natural languages are parasitic on thought. It is just a way to say something that a lot of people have held, that what you do in language is to express thought. If thought has a semantic regimen of some sort or another then all language has to have is a way to express it.

If you look at what seems to be the data about natural languages, they don't look very compositional. If you look back at the tradition of natural-language analysis in philosophy, it seems to me, philosophers assume (and linguists have followed them in this actually) that sentences have very, very abstract representations and are compositional at that level: in philosophy it is logical form and in linguistics it is LF or something. But these are really very abstract. Which is to say that looked at the other way around, you have to do a lot of patching, a lot of fleshing out, as it were, to get from surface English – that is wildly non-compositional – to some underlying structure where you can define entailment rules and so on.

One way to look at it is to say that we discovered that the underlying representations of English are remarkably abstract. Another way of looking at it is to say that these aren't representations of English at all, they are representations of thought. And if you are thinking in something like logical form, then there is no reason to think that the sentences have to have very abstract representations in order to say what you are thinking.

The question now, it seems to me, is whether there is any reason to believe that there is something like a semantic level of *linguistic* representation. I'm inclined to say that maybe there is not. It seems to me, as one laughingly says, a sort of an empirical issue.

I think that what goes on in the actual practice of analysis is that a lot of the motivation for structure assignments in linguistics and philosophy is compositionality. People want to get a representation of sentences so

---

4 See, for instance, Mates (1952)

that they compositionally determine their entailments. But if you consider a picture where you have compositionality of thought and then language just expresses that, it is not clear that those kinds of arguments really have any power. Because they depend on assuming that we got to construct a representational system for language where compositionality is obeyed at some level of abstraction and the abstractions get pretty wild in the literature. If you ask a working linguist how many clauses there are in the deep structures of 'John runs' it turns out that there are about thirty-two or something. And that is a little hard to believe. At least, it is a little hard for me to believe. And I think that that maybe an artifact of the assumption that there must be a level of semantic representation at which *English* is compositional.

MJ: But in the past[5] you established the compositionality of thought partly by way of its systematicity, which in turn was established by way of the systematicity of language. So why is the systematicity of thought indicative of the compositionality of thought, while the systematicity of language is not indicative of the compositionality of language?

JF: Because you can see how language can inherit the systematicity of thought, but there is no way it can be the other way around.

MJ: So the systematicity of *thought* is explained in terms of the compositionality of thought, while the systematicity of *language* is explained in terms of the systematicity of thought**?**

JF: That would be the idea. As I say it is not obvious that you have to do it that way, but it is not obvious that you need to have parallel theories of systematicity as it were. So it is a question for investigation. Who knows?

MJ: The compositionality explanation of the understanding of novel expressions are done in terms of both the meaning of the parts of these expressions and in terms of the syntactical rule by which they have been combined. But among the novel expressions that we actually understand, we find a lot of expressions that are not syntactically well formed. But if you have a compositionality explanation in terms of syntactical rules, that explanation does not seem to apply to these kinds of cases.

JF: I think that is right.

MJ: So what would be an explanation of these cases?

JF: The problem with all this is to define a reasonable theoretical

_____

[5] For instance, Fodor and Pylyshyn (1988)

domain. It looks to me like the well formed sentences are roughly a reasonable theoretical domain. The rest of it is just ways in which we patch up what we know about the language, what we know about the speaker, the situation, and all of that sort of Gricean stuff. I'm not sure you can get a theory of that. Given that you know the language, you can play all sorts of tricks that violate the rules, and I don't think that you can get a theory of tricks.

There are two big things. One is the theory of particular competences like language, and then the other is a theory of what human beings are like as such, as it were. The latter is what Chomsky calls performance kinds of theories. It seems unlikely to me that you can get a theory of language that accounts for everything that people can do with language. There are all sorts of things that some people can do and some people cannot.

MJ: So the reason for construing compositionality in terms of syntactical rules has to do with the need to restrict one's theoretical domain?

JF: Well, the marrow of all scientific investigation is that you have to constrain it somewhere, and you cannot do it a priori. It looks at the moment like something like the well-formed sentences with their truth conditions or something is what it is plausible to have a theory of. Then there are the tricks that intelligent creatures can play with these kinds of structures, and those are just boundless, short of a theory of the mind in general.

IB: What you are saying is that people who argue from communication are simply misunderstanding the domain?

JF: That is the hope. They could be right, in which case we are not doing what we ought to do. I don't think that there is any reason to doubt that the moral of that is to say that these are simply not theoretical domains, we cannot generalize about the world under those descriptions, and the descriptions under which we can generalize about the world turn out to be interesting to a scientist but not to anybody else, roughly speaking.

2. Concepts and categorisation

MJ: Many theories, particularly in cognitive psychology, describe concepts as something essentially related to categorization, which is something that you deny. What is your view of the relation between categorization and concepts?

JF: In a way, the best statement of this that I've seen is in Christopher Peacock's work. The idea is that theories of concepts are parasitic on theories of concept possession –What it is to have a concept, rather than what it is for something to be a concept.

So what is it to have a concept? Well, we have been told for the last fifty years or so that it is to have certain capacities, it is knowing how, not knowing that. What's the knowing how? It is knowing how to sort, or knowing how to infer, or something of that kind. And that is when you get to a dead end. Because you cannot get compositionality out of those.

It is certainly reasonable, I suppose, against a background of a operationalist or a pragmatist theory of content, that the concept CUP is what you use to distinguish between this kind of thing [demonstrates a cup] from that kind of thing [demonstrates a non-cup (a spoon)]. The question is whether that is the individuation conditions of the concepts and it looks like the answer is no. That strikes me as very interesting.

MJ: So what is the relation between categorization and concepts?

JF: You use your conceptual repertoire to effect categorization. Categorization is really a kind of non-demonstrative inference. There is nothing special about the connection between the conceptual repertoire and categorization that doesn't hold, for example, for the relation between the conceptual repertoire and thought. If one is looking for a capacity that is close to a defining capacity for the concept CUP it is the capacity to think about cups. That is what the concept CUP is. It is that mental object that serves to, as it were, bring cups before the mind.

There are some concepts that you don't use to think with. Suppose that Chomsky is exactly right about how language works. Suppose that to understand a sentence, you have to get a phonetic analysis, a syntactic analysis, and God knows what else. What is interesting about that is that the conceptual apparatus that allows you to do that, assuming that the story is true, is not available for you to think with. You can sort noun phrases on the Chomsky kind of story, and that is part of mastering the language, but you cannot think about noun phrases with that apparatus. You will have to go to a linguistics class and learn what noun phrases are and so on.

Here is one way to think about the modularity issue. The fundamental thing you do with concepts is that they are constituents of thought. They allow you to present properties, individuals, and so on, to the mind, but there may be cases where, as a matter of fact, there are architectural restrictions that say 'Look, you can have this conceptual apparatus but you are not allowed to think with it, you are only allowed to do such and such with it.'

It turns out that there is a lot of etiological precedent for that. For instance, it looks like the navigational capacities of very primitive organisms such as insects, are such that what the insects are doing is really very complex mathematics, but they can only do it in the context of locating a food source or something of that sort. They cannot think about triangles, although they can do triangulation very nicely.

3. The conceptual capacities of animals

MJ: Speaking of the cognitive capacities of animals, are there animals that have a language of thought?

JF: You might as well ask: Do animals reason? Do they think? I think that is the way to put it. If they think, the language of thought model seems to be the natural one to apply. The question is whether their intellectual capacities are such that you need that kind of model to explain them, and I don't know. It is an issue of fact.

If you believe in evolutionary explanations, then you expect our intellectual capacities to come from somewhere. You would expect that the representational capacities would increase with phylogeny.

IB: So other primates or at least our closest ancestors the apes would have a language of thought?

JF: Yes. But we are having a terrible trouble figuring out what kind of minds we have, and at least we know we have them. God knows what kind of minds they have. What I would have really thought is this: There has never really been any serious alternative to representational theories of mind. And once you are committed to the notion of mental representation, then there are all sorts of reasons why you would like the representations to be language-like. For example, you need notions like predication to distinguish between thinking about and thinking that. So the basic question is whether or not the cognitive capacitites of animals are such as to require some notion of mental representation. If they are, then it seems very likely that the appropriate notion of mental representation would be language-like.

IB: What do you mean by language-like?

JF: We know that there are some things that we have to have. You have to have reference and you have to have truth roughly speaking. And that means you have to have composition. That is what language-like means. How similar the stuff we or other animals think in is to the stuff in which we communicate in I have no idea. It has to have those basic semantic properties that linguistic representations seem to have, but after that it is an empirical issue, roughly speaking. And I'm not one to prejudge it a priori. We have to attribute to them whatever we need in order to explain their capacities.

MJ: But given that they have a language of thought, how come they don't have a language – if our language of thought is what enables us to have a language?

JF: That is a good question. I'm looking at it from the other way. What the language of thought is for, is for mental representation. It is for thinking with. It is the medium that you think in. What is clever about our species is that we found a way to communicate with something like the system that we think with. On this kind of view that was just some sort of biological jump, something that happened to us, but didn't happen to other kinds of animals. What you would expect from a language of thought-story is not that other kinds of animals could talk it is that other animals could think. And that, it seems to me, there is no reason to deny.

4. Philosophy today

MJ: Considering that you have been working within philosophy now for over forty years and have had a great impact on the discipline, can you, in closing, tell us about your views on the role philosophy plays today, or for that matter, the role that you think *analytical* philosophy plays today?

JF: I don't think there is any such thing as analytical philosophy. There is the question of whether philosophers should give arguments. It is a good idea to give arguments. Not because anybody is ever persuaded by them, but because they allow you to crystallize and encapsulate the sense of a geography of an issue. I really don't think there is such a thing as conceptual analysis. I don't think that concepts have analyses. Quine was right about that. And once one has seen the point of the Quinean polemics, there just is no such process as analyzing a concept. That is why we have analyses for so few of them. So if that is what analytic philosophy is,

trying to give definitions, or definitions in use, or criteria, or something of that sort, I think that is a bankrupt enterprise.

What's interesting in conceptual analysis, what's interesting for me anyway, is figuring out what the geography of various kinds of theories is. That is something that people in psychology, for instance, are, by and large, not well-trained to do. I mean they are not encouraged to step back from a problem and ask: 'What kind of theory am I trying to apply here? To what kind of problem am I trying to apply it?' They become extremely impatient when you ask these kinds of questions. And philosophers are trained to do that. That is as close to conceptual analysis as I get.

MJ: So you see an important role for philosophers to play in other disciplines?

F: On this kind of view, yes. Especially where you have relatively unformulated disciplines, relatively new disciplines. The practitioners of these disciplines tend to be very badly trained from the point of view of assessment and construction of theoretical mechanisms.

Here I'm quite sympathetic to the intuitions, but not the practice, of people like Wittgenstein, that is, one has the sense of not knowing one's way around, one has the sense of not understanding the apparatus one is using in defining one's research project, giving explanations, and so on. That is the kind of thing that philosophical training sometimes helps with. So I don't think that the fate of the discipline depends upon being able to give an account of what a philosophical analysis is in the specialized sense people have had in mind.

And in fact, for the last two thousand years or so, that is not mostly what philosophers have done. Mostly what philosophers have done is construct theories. When you think that the concepts are a mess, you try to straighten them out a bit and when you don't, you leave them in situ. It seems to me that what we are doing is vastly more like what Kant, or for that matter Hume, thought he was doing rather than what Ryle or Wittgenstein thought that they were doing. The analytic stuff, strictly so called, was just a methodological aberration which we are now getting over.

References

FODOR, J.A. (2001). "Language, Thought and Compositionality." *Mind and Language* **16**(1): 1–15.

FODOR, J.A. and E. LEPORE, Eds. (2002). *The Compositionality Papers*, Oxford University Press.

FODOR, J.A. and Z. PYLYSHYN (1988). "Connectionism and cognitive architecture: A critical analysis." *Cognition* **28**: 3–71.

FREGE, G. (1892). On Sense and Meaning.

HODGES, W. (2001). "Formal Features of Compositionality." *Journal of Logic, Language, and Information* **10**(1): 7–28.

MATES, B. (1952). Synonymity. *Semantics and the Philosophy of Language*. L. Linsky. Urbana, IL., University of Illinois Press.

A New Definition of Endurance[*]

by

MILLER

Department of philosophy, University of Queensland

Abstract In this paper I present a new definition of endurance. I argue that the three-dimensionalist ought to adopt a different understanding from the four-dimensionalist, of what it is to have a part *simpliciter*. With this new understanding it becomes possible to define endurance in a manner that both preserves the central endurantist intuitions, whilst avoiding commitment to any controversial metaphysical theses. Furthermore, since this endurantist definition is a mereological one, there is an elegant symmetry between the definitions of endurance and perdurance in that the theories of three- and four-dimensionalism are both expressed in the language of mereology. Nevertheless, though both definitions are expressed within the same broad language, some of the terms of that language have subtly different meanings within the context of each theory. It is in understanding on the one hand that each theory is essentially a mereological theory and that therefore each shares some underlying theoretical similarities, and yet also that there are some subtle differences in the way each theory understands some of the terms of mereology, that allows us to see clearly what lies at the heart of the debate between these two accounts of persistence.

Keywords: endurance, persistence, three-dimensionalism, four-dimensionalism, identity

1 Introduction

FOR SOME TIME now, there have been two main competitors when it comes to accounts of persistence: three- and four-dimensionalism. Three-dimensionalism is the thesis that persisting objects have only spatial dimensions. Thus since our world is a world in which there are three spatial dimensions, in our world persisting objects are three-dimensional. Four-dimensionalism, on the other hand, is the view that persisting objects have both spatial dimensions and a temporal dimension. Since our world is a world with three spatial dimensions and one temporal dimension, this is the view that persisting objects are four-dimensional: they are extended in time as well as in space.

The three-dimensionalist holds that all objects are wholly present

---